

# ADAPTIVE SOFTWARE MANAGEMENT

## BACKGROUND OF THE INVENTION

### Field of the Invention

The invention relates to adaptive control of electronic, electromechanical and mechanical devices. More particularly, it relates to techniques for managing devices that  
5 are controlled using computer software.

### Description of the Related Art

Computer controlled systems and devices are ubiquitous in today's society and there has been a long-felt need to optimize their operation to adapt to the needs of the individuals who use them. One such class of devices in which adaptive control is needed  
10 is portable battery powered devices. With the proliferation of battery powered devices, power management has become critical to the success of such products. Power drain effects all aspects of operation: size, weight, battery life. Adequately managing power allows for use of smaller and lighter batteries, other advanced power technologies (e.g., fuel cells, etc.), and other alternative portable power sources, each of which can support  
15 longer periods of operation of the devices. Although power management can be accomplished according to many approaches, none are known that are adaptive to a user's patterns of use or robust enough to optimize the power management in the presence of noise.

Designers of today's electronic devices attempt to reduce power consumption by  
20 concentrating their efforts on reducing the power drain required by the device or by using background processing tasks to determine if certain subsystems within the electronic

device can be powered down. These devices often are described as being adaptive, since they feed back to a power management algorithm information concerning the power usage of the device to determine new values for adjustable power settings. However, the degree of adaptability of those device is very low since they employ only a mere feed back loop  
5 that does not account for system noise or provide control in dynamic situations.

Accordingly, there is a need for power management techniques that take into account system noise, user preferences, and patterns of use of the device. Power management systems are but one class of systems that are in need of adaptive control techniques. Many other types of systems can benefit from adaptive control techniques.

## 10 SUMMARY OF THE INVENTION

Therefore, in light of the above, and for other reasons that will become apparent when the invention is fully described, an aspect of the invention relates to a method of generating information for adaptively controlling a device.

For example, the power consumption of a device can be adaptively controlled  
15 according to the methods described here. One such method includes performing an optimization analysis based on a quality of service and power consumption of the device, and a plurality of control factors for controlling the device. A power management profile is generated that relates to the quality of service and power consumption of the device based on results of the optimization experiment. The power management profile is used to  
20 automatically determine a control factor to adjust, and a level to adjust it to, to minimize power consumption or achieve a certain quality of service level.

The above features and advantages of the invention will become apparent upon consideration of the following descriptions and descriptive figures of specific

embodiments thereof. While these descriptions go into specific details of the invention, it should be understood that variations may and do exist and would be apparent to those skilled in the art based on the descriptions herein.

## BRIEF DESCRIPTION OF THE DRAWINGS

5           Fig. 1A is a block diagram illustrating a power management system using off-line adaptive software power management techniques described here.

          Fig. 1B is a block diagram illustrating a power management system using both off-line and on-line adaptive software power management techniques described here.

          Fig. 2A is a graph showing the effects on a power consumption factor of a device,  
10   measured as a signal to noise ratio (SNR), as two control factors for that device are varied among different levels. The points on the graph are generated from experiments run on a system that uses a digital signal processor (DSP) as the object under test that operates in various modes and charting the resulting power consumption of the DSP measured as a signal to a noise ratio (SNR).

15           Fig. 2B is a graph showing the effects on a quality of service (QoS) attribute of a device, here, the wake-up latency of the DSP which has several states of suspended operation to conserve power, as the two control factors for the DSP are varied among different levels. The points on the graph are generated from experiments run on the DSP operating in various modes and charting the resulting wake-up latency of the DSP  
20   measured in milliseconds.

          Fig. 3A is a graph illustrating a relation between a quality of service measure (i.e., wake-up latency) and power consumption of the DSP, generated based on the graphs of Figs. 2A and 2B. The line connecting points 1, 2, 4, 5, 7, 8 and 9 indicates the optimal

operating conditions that satisfy a QoS metric, in this case wake-up latency, while minimizing power consumption.

Fig. 3B is another view of the graph in Fig. 3A, illustrating a relation between power consumption of the DSP and the quality of service measure (i.e., wake-up latency),  
5 generated based on the graphs of Fig. 2A and 2B.

Fig. 3C is a graph illustrating example of another quality of service measure, namely, the brightness of a back light in a liquid crystal display (LCD) in relation to power consumption of the device employing the LCD.

Fig. 3D is another view of the graph in Fig. 3C illustrating a relation between  
10 brightness of an LCD back light and power consumption of the device.

Fig. 4 is a block diagram showing the off-line adaptive power management system of Fig. 1A, in greater detail.

Figs. 5A and 5B are block diagrams showing the off-line and on-line adaptive power management system of Fig. 1B, in greater detail.

15 Fig. 6 shows a block diagram of a hardware system for implementing the adaptive power management systems shown in Figs. 4, 5A and 5B.

Fig. 7 is a block diagram showing a power management processor.

## DETAILED DESCRIPTION

The embodiments described below are described with reference to the above  
20 drawings, in which like reference numerals designate like components.

The adaptive control techniques described here can be used in a wide variety of computer-controlled systems and devices. Virtually any system or device that is controlled

by a computer and that is suitable for operating according to an individual operator's preferences or habits can benefit from the adaptive control techniques described here. For example, the adaptive control techniques described here can be applied to HVAC control systems, lighting control systems, home automation system, etc.

5           To illustrate the inventive techniques, embodiments that relate to power management are described here. For example, methods, devices and articles of manufacture are described here that can dramatically improve power management of electronic devices by optimizing the performance of the device in the presence of compound system noise. By using experiments, such as conventional Design of  
10 Experiments (DOE) experiments, Orthogonal Array experiments, Latin Square Design experiments, Taguchi style design of experiments, and full and fractional factorial style experiments, control factors can be chosen that control the device so that it performs in an optimal manner based on those chosen control factors, even in the presence of system noise. These experiments can be performed in an off-line mode to characterize the  
15 device's response for use when operating the device. Certain quality of service (QoS) performance metrics are selected to gauge the effect of changing a control factor, or design constraint. Anticipated use conditions, representing noise in the experiment, are factored into the experiments. The experiments are run resulting in a compilation of data that represents the effect on the QoS parameters as the control factors are varied, as well as the  
20 effect on power consumption. This data is analyzed to produce subsequent data typically in the form of plots and tables that describe how QoS varies as power consumption is varied, and vice versa.

A power management unit in the electronic device uses these QoS/power consumption plots during operation to manage the device's power consumption. The

power management unit operates according to a power management strategy that is either predetermined or designated by a user. The power management unit monitors the relevant QoS parameters and the power consumption status. Based on the QoS/power consumption plots, generated in an off-line mode, the power management unit employs the selected  
5 strategy to vary a control factor to achieve a desired QoS or power consumption.

Alternatively, the QoS/power consumption plots can be updated in an on-line mode by monitoring the real-time QoS and power consumption, and performing the experiments and subsequent data analysis while the device is operating based on the measured QoS and power consumption data. In this manner, an on-line control of the device is achieved.

10 Off-line power management typically is performed at either the design or manufacturing stages of development to provide a product with “out of the box” power management. For some products, such as pagers, non-configurable items, or items used by different people, this is the only practical type of power management available. Other products that have processing capability can be customized by the user using an on-line  
15 power management technique. On-line power management is performed continuously in the background, adjusting various controllable settings to minimize the power usage of the product.

Off-line power management is appropriate to determine the initial/factory settings for power optimization for all products. For non user-configurable products, such as  
20 pagers or fire safety radios, where the user obtains equipment from a pool of radios, the off-line parameters would be used as the default power management settings and would be continuously monitored to provide either the best quality of service (QoS) or maximum battery life, depending upon the application.

In off-line power management, a series of analysis are performed that characterize the relationship between the static design parameters and desired output characteristics (e.g. QoS, power consumption, etc.). For battery powered devices these design parameters might include processor type, battery characteristics, RF settings, and option selection  
5 (back light, beep/vibrate, auto power down, etc). Output characteristics include, but are not limited to battery life, quality of service, brightness, and other features available to the user.

Once the optimal design settings are determined they provided a fixed, but minimized, power profile.

10 On-line power management provides for continuous monitoring, optimization and updating of the design settings to minimize the in-use power consumption. Optimization of these settings is based on previously identified factory criteria, user preferences and actual in-use conditions. For example, in a personal digital assistant (PDA) design, the off-line experiments might determine that a light sensor has a significant effect on the back light  
15 power consumption. The online system would monitor not only this static level relationship but also would allow for the user to specify the amount of back light requirement for his or her particular comfort level. Thus, the user can choose the best quality of service for his or her particular environment. Having specified a particular user constraint, the system can vary other parameters to ensure the appropriate level of service.  
20 Various features can be disabled based on the user selections.

As another example, a laptop computer writing to disk consumes a great deal of power. The on-line power management technique monitors the battery usage and changes the caching technique to reduce disk access. Although this raises the risk of lost data, the user can select his or her level of risk, based on the criticality of information being stored.

In yet another example, military, police or fire protection radio users who are issued a radio for long term use typically want to set custom features that are unique for their application. Such users want to have pre-set frequencies and menus to reduce the time that the display is operational. How and when the display is powered down is a factor  
5 in the on-line power management. The user may require the device to remain powered on for an eight (8) hour operational mission, and that requirement would be a major factor in determining the QoS that can be delivered. For instance, the radio user might select a typical percentage of expected voice, video and data traffic. The on-line monitoring can alert the user if the pattern of use does not agree with the request, and optionally make run-  
10 time decisions to meet the required eight (8) hour mission. Thus, instead of running out of power unexpectedly, the device can adjust the QoS or alert the user as appropriate.

Cell phone users typically want to set custom features that are unique the user's preferred style of using the cell phone. For example, cell phone users typically want to set the update frequency for the voice mail or roaming ping messages to locate a nearby  
15 transmission tower. Using the power management techniques described here, the user can choose a longer acquisition time to provide longer battery life.

Fig. 1A is a block diagram of an adaptive power management system for an electronic device, operating in an off-line mode to establish power management strategies that can be used to control the power consumed by the electric device. The system shown  
20 in Fig. 1A operates prior to using the electronic device, to establish in advance the power management strategies to be used in powering the device. In this manner, the power management strategies are generated off-line prior to using the device.

The off-line adaptive power management system includes an off-line power management strategy analysis unit 10. The analysis unit 10 includes software that



performs off-line experiments to characterize various design constraints and QoS measures of the device. Inputs to the experiments include anticipated use conditions, QoS performance metrics, and various design constraints (or control factors) that are device dependent. By performing the off-line experiments, power management profiles are developed. An off-line QoS/power management profile unit 20 receives the results of the experiments and generates profiles based on those results. The generated power management profiles are stores in the off-line profile unit 20. The power management profiles held in the profile unit 20 are input to a run-time control unit 30 that controls the device based on the device's operating conditions and the power management profiles.

Fig. 1B is a block diagram of an off-line/on-line power management system. Here, an off-line power management strategy analysis unit 10 receives inputs that are similar to the inputs to the off-line power management strategy analysis unit 10 shown in Fig. 1A, namely, anticipated use conditions, QoS performance metrics, and design constraints. The output of the analysis unit 10 is provided to an off-line QoS/power management profile unit 20 that generates and stores power management profiles. Those power management profiles are provided to a run-time control unit 30 that operates and controls the device based on the power management profiles and operating conditions.

The off-line/on-line power management system shown in Fig. 1B also includes an on-line section. The on-line section includes an on-line analysis unit 40 and a user profile/usage history unit 50. The on-line analysis unit 40 operates to analyze in real time various constraints and parameters of the device based both on the real time operating conditions and based on user profile information or user history information from unit 50. The on-line analysis unit 40 outputs to the profile unit 20 information to update the power management profiles that are used by the run time control unit 30. This forms a control

loop that adapts the power consumption to the actual operating conditions of the device and the user's patterns of use of the device.

Fig. 2A is a power-consumption/control factor effects graph. This graph plots certain control factors that affect the device's power consumption with respect to a QoS metric. These control factors can include different design constraints of the device. Fig. 2A illustrates two such control factors labeled "A" and "B." A first graph 60 shows the change in power consumption, measured as a signal-to-noise ratio in decibels (dB), of the device based on different levels for the control factor. In this case, factor A is varied among three levels, A1, A2 and A3. The design experiments are run by varying the control factor levels and measuring the power consumption for each level. For example, when factor A is set at a first level (A1), the device QoS measurement, namely a signal-to-noise ratio (SNR), it is measured at approximately -36dB. When factor A is set to a second level (A2), the measured signal-to-noise ratio is approximately -35dB. At the third level for factor A (A3) the signal-to-noise ratio is measured at approximately -32dB.

An example of a control factor, labeled A in Fig. 2A, is an idle mode instruction in a digital signal processor (DSP) integrated circuit (IC). For example, a DSP manufactured by Texas Instruments, e.g., DSP model no. TMS320LC54x, has three power-down modes that can be set to control various states of dormancy for the processor. Each state disables certain functions of the DSP to reduce, by varying degrees, the power consumption of the DSP. By entering one of these power-down modes, the DSP will consume less than the amount of power consumed during normal operations. These power-down modes are often referred to as suspend modes or sleep modes. In this example, the DSP can be set to enter various dormant states by issuing, under software control, the instructions: IDLE\_1,

IDLE\_2, or IDLE\_3 instructions. Table 1 below shows portions of the DSP that are affected by these software commands.

Table 1: Low Power Operational Modes

<b>Operation</b>	<b>IDLE_1</b>	<b>IDLE_2</b>	<b>IDLE_3</b>
CPU Halted	Yes	Yes	Yes
Peripheral Clock Halted	No	Yes	Yes
PLL Halted	No	No	Yes
Data/Address State	No	No	No
External Control State	No	No	No
Internal HW Interrupts	No	Yes	Yes
External HW Interrupts	No	No	No

5 Table 2 below shows two control factors: the IDLE mode instruction, labeled control factor “A,” and the phase-lock loop (PLL)/clock rate in the DSP, labeled control factor “B.” Table 2 shows three different levels to which each control factor can be set.

Table 2: Control Factor Levels

<b>Control Factor</b>	<b>Factor Label</b>	<b>Level 1</b>	<b>Level 2</b>	<b>Level 3</b>
Idle Mode Instruction	A	IDLE_1	IDLE_2	IDLE_3
PLL/Clock Rate (MHz)	B	100	50	25

10 As illustrated in Fig. 2A, if control factor A1 corresponds to the IDLE\_1 instruction, A2 corresponds to the IDLE\_2 instruction, and A3 equals corresponds to the IDLE\_3 instruction, then Fig. 2A can be shows varying degrees of power consumption, measured as a signal-to-noise ratio in dB. For example, with control factor A1 selected

(IDLE\_1 instruction) the power consumption is less than when control factor A3 is selected (IDLE\_3 instruction). In this manner, by performing experiments on the device with the control factors set to various levels, the plot shown in Fig. 2A is generated. This plot characterizes the quality of service of a device in relation to the various control factor  
5 levels.

Also shown in Fig. 2A is a second control factor, control factor B, set to various levels. Here, when control factor B1 is set, the device exhibits a power consumption, measured as a signal-to-noise ratio of approximately -38dB. When control factor B2 is set, the power consumption is approximately -34dB. When control factor B3 is set, the power  
10 consumption raises to approximately -31dB. An example of a DSP control factor corresponding to control factor B is a frequency of a phase lock loop (PLL) circuit. That is, the PLL clock frequency, or speed, can be set to three different values in the DSP. That is, the PLL clock can be set to 100 MHz, 50 MHz, or 25 MHz.

A quality of service measure for this DSP example is a wake-up latency time of the  
15 DSP. The wake-up latency time is a measure of the amount of time required between commanding the DSP to exit an idle mode and the DSP being ready to operate in its normal mode of operation.

Table 3 below is a compilation of an exhaustive sequence of n-way experiments, varying each of the idle mode and PLL clock frequency control factors.

Table 3: Tabulated Results of Experiments that Vary Each of the Control Factors.

<b>Experiment No.</b>	<b>Idle Mode</b>	<b>PLL Clock (MHz)</b>
1	IDLE_1	100
2	IDLE_2	100
3	IDLE_3	100
4	IDLE_1	50
5	IDLE_2	50
6	IDLE_3	50
7	IDLE_1	25
8	IDLE_2	25
9	IDLE_3	25

Table 4 shows the results of each of the experiments identified in Table 3.

Table 4: Experimental Results

Experiment No.	Power SNR (dB)	Wake-up Latency (ms)
1	-40.2	0.00
2	-40.1	0.00
3	-35.0	0.14
4	-35.7	0.00
5	-34.5	0.00
6	-31.8	0.09
7	-32.4	0.00
8	-31.2	0.00
9	-29.7	0.07

The results from the experiments shown above in Table 4 are plotted in the graphs of Figs. 3A and 3B. In Fig. 3A the graph 60 shows how power consumption changes when the idle mode instruction changes. Graph 70 shows how power consumption changes as the PLL clock frequency changes. In Fig. 2B, a graph 80 shows how the DSP wake-up latency changes as the idle mode control factor changes. Graph 90 shows how the DSP wake-up latency changes as the PLL clock frequency control factor changes.

Fig. 2B illustrates another measure of quality of service, namely, the wake up latency of the DSP. The wake up latency is measured in milliseconds (ms) and corresponds to the amount of time the DSP takes to exit a background thread to process an interrupt. Here, when control factor A1 is set there is virtually no latency. Similarly, when control factor A2 is set there is virtually no wake up latency. However, when control factor A3 is set a wake up latency of 0.10 milliseconds is experienced. Varying control factor B produces the following results: control factor B1 causes a wake-up latency of

approximately 0.05 milliseconds; control factor B2 causes a wake-up latency of 0.03 milliseconds; and control factor B3 causes a wake-up latency of approximately 0.02 ms.

Figs. 3A and 3B show a plot of the power consumption versus quality of service measures based on the experiments illustrated in Figs. 2A and 2B. Fig. 3A shows a plot that relates the quality of service to the power consumption. Fig. 3B shows a plot that is the inverse of the plot shown in Fig. 3A. Namely, it illustrates a relation between power consumption and quality of service.

The QoS/power consumption charts produced from the experiments can be filtered to remove extraneous data points from the plots. This is illustrated in Figs. 3A and 3B in which a curve is fitted to the data such that the data points corresponding to experiment numbers 3 and 6 are filtered out of the relation.

Another example of a QoS measure related to power consumption is the brightness of a display back light. In portable electronic devices that employ an liquid crystal display (LCD), a back light is often needed to enhance the readability of the display. However, using the back light can significantly decrease the life of batteries powering the device. Accordingly, users are often faced with a tradeoff in using the back light at the expense of depleting the charge on batteries that power the device. The adaptive power control techniques described here help manage the power in such a device having a back lit LCD display, while minimizing the drain on batteries powering the device.

Fig. 3C shows a curve fit to data points that would be experimentally produced that relate power consumption to the brightness of the back light display, measured in candles per square meter. Here, the graph illustrates that as brightness of the back light increases, the power consumption increases, almost exponentially. In this example, if the brightness

control factor is desired to be  $150 \text{ cd/m}^2$ , the expected power consumption will be 500 mW.

Similarly, the graph in Fig. 3D plots the same data points from the experiment, but relates back light brightness to the amount of power consumption. If a power consumption of 500 mW is desired, the resulting back light brightness will be  $150 \text{ cd/ m}^2$ .

A block diagram of an offline adaptive power management system is shown in Fig. 4. The system is divided into two portions. The first is a design stage that relates to designing the power management strategies to be used in the device. It includes performing the offline power consumption/QoS optimization experiments and generating the QoS and power consumption graphs or tables for use in adaptively controlling the device.

The second portion of the device includes the run-time processing components. This portion operates with the device to monitor the power consumption and certain QoS metrics. That information is used together with the power consumption/QoS graphs generated offline to control the device based upon criteria a user designates as important, such as maximizing operation time, or ensuring a certain quality of service such as maintaining a certain degree of back light brightness.

Referring to Fig. 4, the design stage 180 includes performing optimization experiments 200 to optimize power consumption and quality of service measures. Examples of the type of optimization experiments are DOE experiments, Orthogonal Array experiments, Latin Square Design experiments, Taguchi style experiments, and full and fractional factorial style experiments, although it will be understood that other types of optimization experiments can be used. Prior to performing the experiments a power strategy designer must quantify static design constraints 210. At this stage, a designer



determines the design constraints that can be used as control factors to the control the device to effect the appropriate quality of service. For example, design constraints, or control factors, that affect both power consumption and a quality of service parameter of interest are selected. In the example of the DSP, the design constraints that are chosen  
5 include the DSP's idle mode instructions and adjustable PLL clock speeds. Here, engineering judgement is exercised to select the most appropriate design constraints for the particular device to be controlled.

At stage 220, the designer specifies and quantifies anticipated user conditions. Here, the designer specifies how target users most likely will use the device in practice.  
10 For example, if the intended users primarily will be operating the device in a low-light environment, the designer can specify an anticipated use condition as one of a low-light environment.

In stage 230, the designer specifies and quantifies the appropriate quality of service and power performance metrics. Here, the designer identifies and characterizes the quality  
15 of service and power metrics according to an application of the device. For example, if the device is a radio for use in a military field operation, a long power life would be designated as having a high priority to ensure the radio can be powered while in the field for an extended period of time. However, the voice quality need only be set to have a discernable quality. Another application of the same radio might have different QoS and  
20 power performance metrics. For example, if the device is used in a firefighter's radio, long battery life extending over several days might not be as important as high voice quality.

These three inputs are supplied to the power/QoS optimization experiments in stage 200. Here, any of a variety of well-known optimization experiments can be employed, such as Taguchi experiments; fractional factorial experiments; full factorial experiments;

DOE experiments, Orthogonal Array experiments, Latin Square Design experiments, etc.

The optimization experiments generate experimental data upon which a variance analysis

240 is performed, for use in characterizing the data. In stage 250 n-way plots and/or tables

are generated of the relation between the control factors and the QoS metrics, similar to the

5 plot shown in Fig. 2B. Similarly, in stage 260, the resulting data from the experiments is

used to produce an n-way plot or table of the relation between the control factors and

power consumption, similar to the plot shown in Fig. 2A.

Stage 270 generates the QoS versus power consumption plot, and the power

consumption versus QoS plots, such as the plots shown in Figs. 3A and 3B, respectively,

10 and Figs. 3C and 3D, respectively. At stage 270 the experimental data is also filtered to

remove any outlying data points. A smooth curve is fit to the data to characterize the

power consumption versus QoS plot, which is made available in stage 290. Similarly, a

smooth curve is fit to the data to characterize the QoS versus the power consumption plot

and is made available in stage 300.

15 Based on the QoS and power consumption data generated in stage 270, default

values of the power management design constraints are generated in stage 280. These

default design constraints, or control factors, are set for use with the device.

The run-time processing portion 190 of the system is illustrated in Fig. 4. The

outputs from the design stage, namely, the default control factor values, and the power and

20 QoS plots and/or tables, are input to a power management unit for the device. The device,

in stage 310, monitors the device's power consumption and QoS metrics. Here, the power

management control factors, or design constraints, are adjusted in stage 320 based on the

monitored power consumption and QoS metrics, and based on the power consumption

versus QoS plot/table from stage 290 and on the QoS versus power consumption plot/table

300. The adjusted control factors are then used to control the device in stage 330. The adjustment/control stages continue during operation of the device.

An on-line power management system is shown in Figs. 5A and 5B. Here, much of the design stage and run time processing portions of the system are the same as with the off-line system shown in Fig. 4. However, in addition to those off-line features, features shown in Fig. 5B are added. Optimization experiments are performed in stage 360 based on the monitored power consumption and QoS metrics from stage 310. Also input to the on-line power/QoS optimization experiment stage 360 are user profiles and usage history information that is collected and stored at stage 370, and other actual usage conditions monitored at stage 380. The optimization experiments are performed, preferably at periodic intervals appropriate to the device or design. Although, the experiments can be performed on an event basis.

As with the off-line experiments, the experiment data is used to perform a variance analysis in stage 390. Also, n-way control factor effects versus QoS performance plots and/or tables are prepared at stage 400, and n-way control factor effects versus power consumption plots and/or tables are prepared at stage 410.

The variance analysis information, and plots and/or tables generated at stages 400 and 410 are supplied to stage 420. There, the QoS versus power consumption and the power consumption versus QoS curves are updated with the newly generated optimization data based on the actual operating conditions.

These updated plots and/or tables are supplied back to stages 290 and 300 where the updated plots and/or tables are used to adjust the power management control factors at stage 350.

A block diagram of a device using the power management techniques described here is shown in Fig. 6. The device includes a sensor interface 430 that receives information from hardware sensors in the device. A power control interface 440 is connected to the various subsystems within the device and operates to output control signals to those subsystem to effect the adaptive control. The sensor interface and power control interface are coupled to a power management application programming interface (API) 450. This unit provides a programming interface to a performance table storage unit 460 that holds the power consumption/QoS plots. The power management API 450 also is coupled to an on-line power optimization API 470.

The on-line power optimization API 470 provides a programming interface to various hardware analysis plug-in units for analyzing the experimental data. An experimentation plug-in unit 490 allows various conventional optimization routines to be used with the system. The on-line power optimization API takes the experiment and variance information and updates the plots and tables in the performance table 460. That updated table is used by the power management API to control the hardware through the power control interface 440 to effect the desired power management.

The power management API 450 and the power optimization API 470 include software, or program instructions, that when executed by a computer perform the adaptive management functions described here. It will be understood that those program instructions can be recorded on a variety of computer-readable media, including, but not limited to magnetic storage media, optical storage media, and electronic storage media. It will be further understood that those program instructions can be transmitted over data networks such as the Internet. Likewise, the performance table storage unit 460 can include any storage media suitable for storing and retrieving the power consumption/OoS

plots, such as magnetic, optical and electronic storage media. Further, such plots can be transmitted over data networks such as the Internet.

Fig. 7 illustrates the TMS320LC54x that can be used as a power management DSP. The power management DSP receives control factor instructions, such as the idle instruction controls and instructions regarding the PLL clock rate. Noise factors, such as program activity and bus activity within the device are monitored and supplied to the DSP. The DSP performs the experiments using the input control factors and the noise factors, and produces the results of the experiments, such as the power consumption and the wake-up latency.

Set forth below are some additional power management/control applications that can use the adaptive control techniques described here. Although the example applications described below relate mainly to mechanical/electromechanical applications, the techniques described here can be applied to virtually any computer controlled system or device that can be operated based on an individual's preferences.

#### Automobiles

The adaptive control techniques described here can be used to control systems in automobiles. For example, engine fuel efficiency can be monitored and settings for the automobile can be controlled using driver conditions and preferences, including the length of trip, remaining fuel, passenger comfort (e.g. heating/air conditioning), etc.

#### Home and office heating and cooling/HVAC systems

Through the use of various room sensors an HVAC system using the techniques described here can develop a optimal energy efficiency profile with respect to a QoS measure such as a desired comfort level. By combining measurements from a dispersed set of motion and temperature sensors the HVAC controller obtains the necessary

information to optimize energy usage. This control also can be based upon other factors, such as the air humidity and the anticipated power loading to the power grid. Users who want to save money might choose to sacrifice comfort in exchange for cost savings during peak usage hours. By using an HVAC system controlled according to the techniques  
5 described here those tradeoffs can be accomplished in an automatic manner.

### Lighting systems

A lighting system can use motion sensors, such as those used in connection with the HVAC system described above, to provide an optimal lighting scenario based on individual desires. The lighting system, using the adaptive control techniques described  
10 here, adapts control of certain lights in the system by controlling the desired levels of light output for those lights with respect to, for example, time of day and room usage by the occupants.

### Airline jet engines

An adaptive power management unit can use existing airframe sensors to optimize  
15 an airline engine's fuel efficiency. Based on speed, environment, altitude, range to touch down, the engine control system throttles and controls the engine using the techniques described here to maximize the QoS, such as the range to touch down, for example. Alternatively, if the QoS measure is passenger comfort, as might be expected for a corporate jet, the adaptive control techniques might alter altitude and speed to improve the  
20 quality of the ride.

Also, the jet engines can be controlled over heavily populated areas to minimize the noise pollution while maintaining the touch down time and passenger comfort. The controller would have information about the airplane's location due to onboard GPS equipment, for example, and would use software executing on a computerized controller to

automatically minimize jet engine noise as it approached heavily populated areas. For instance, the Concorde can be controlled so that it does not reach supersonic speeds until it is located over the ocean due to the sonic booms that it creates as a result of its super-sonic speed.

5    Motors, pumps, etc

          Motors and pumps can be optimized to control flow with respect to energy draw. Various pump parameters can be controlled using the techniques described here to maintain laminar flow, continuous ramp up, and consistent volume. The advantage of such a system is that it is adaptable to the users' needs as opposed to being fixed for all users  
10   irrespective of any specific user's QoS requirements. Such a motor control system also can be adapted to other factors, such as rain flow. Over a heavy period of rain, it might be more important to pump the water to prevent flooding, regardless of the energy costs. The motor control system using the adaptive control techniques described here can be configured to adapt to those types of conditions.

15           Having described embodiments of devices, articles of manufacture and methods for adaptively controlling power management in devices, it is believed that other modifications, variations and changes will be suggested to those skilled in the art in view of the teachings set forth herein. It is therefore to be understood that all such variations, modifications and changes are believed to fall within the scope of the present invention as  
20   defined by the appended claims. Although specific terms are employed herein, they are used in their ordinary and accustomed manner only, unless expressly defined differently herein, and not for purposes of limitation.